

Erlang 4.4 Extensions

EDE Team

1997-05-21

Typeset in \LaTeX from SGML source using the DOCBUILDER 3.0 Document System.

Contents

1 Erlang 4.4 Extensions	1
1.1 Records	2
Defining a Record	2
Including a Record Definition	3
Creating a Record	3
Selectors	4
Updating a Record	4
Type Testing	5
Pattern Matching	5
Nested Records	6
Internal Representation of Records	6
Example	7
1.2 Programming with Funs	9
Higher Order Functions	9
Advantages of Higher Order Functions	11
The Syntax of Funs	11
Variable Bindings within a Fun	12
Funs and the Module lists	13
Funs which Return Funs	18
1.3 List Comprehensions	21
Examples of List Comprehensions	21
Variable Bindings in List Comprehensions	23
1.4 Macros	25
Macros and Tokens	25
Pre-Defined Macros	26
Flow Control in Macros	26
A Macro Expansion Utility	27

Chapter 1

Erlang 4.4 Extensions

This chapter describes the Erlang 4.4 extension. The chapter contains the following sections:

- Records
- Functional Objects (Funs)
- List Comprehensions
- Macros.

1.1 Records

A record is a data structure intended for storing a fixed number of related data items. It is similar to a struct in C, or a record in Pascal.

The main advantage of using records instead of tuples is that fields in a record are accessed by name, whereas fields in a tuple are accessed by position. To illustrate these differences, suppose that we want to represent a person with the *tuple* {Name, Address, Phone}.

We must remember that the Name field is the first element of the tuple, the Address field is the second element, and so on, in order to write functions which manipulate this data. For example, to extract data from a variable P which contains such a tuple we might write the following code and then use pattern matching to extract the relevant fields.

```
Name = element(1, P),
Address = element(2, P),
...
```

Code like this is difficult to read and understand and errors occur if we get the numbering of the elements in the tuple wrong. If we change the data representation by re-ordering the fields, or by adding or removing a field, then all references to the person tuple, wherever they occur, must be checked and possibly modified.

Records allow us to refer to the fields by name and not position. We use a record instead of a tuple to store the data. If we write a record definition of the type shown below, we can then refer to the fields of the record by name.

```
-record(person, {name, phone, address}).
```

For example, if P is now a variable whose value is a person record, we can code as follows in order to access the name and address fields of the records.

```
Name = P#person.name,
Address = P#person.address,
...
```

In the following sections we describe the different operations which can be performed on records:

Defining a Record

A record is defined with the following syntax:

```
-record(RecordName, {Field1 [= DefaultValue1],
                    Field2 [= DefaultValue2],
                    ...,
                    FieldN [= DefaultValueN]}).
```

The record name and field names must be atoms. The optional default values, which are terms, are used if no value is supplied for a field when a new instance of the record is created. If the default value is not supplied, then the atom `undefined` is assumed.

For example, in the following record definition, the `address` field is `undefined`.

```
-record(person, {name = "", phone = [], address}).
```

This definition of a person will be used in many of the examples which follow.

Including a Record Definition

If the record is used in several modules, its definition should be placed in a `.hrl` header file. Each module which uses the record definition should have a `-include(FileName)` statement. For example:

```
-include("my_data_structures.hrl").
```

Note:

The definition of the record must come before it is used.

Creating a Record

A new record is created with the following syntax:

```
#RecordName{Field1=Expr1,  
            ...,  
            FieldM=ExprM}.
```

If any of the fields is omitted, then the default value supplied in the record definition is used. For example:

```
> #person{phone = [0,8,2,3,4,3,1,2], name = "Robert"}.  
{person, "Robert", [0,8,2,3,4,3,1,2], undefined}.
```

Selectors

The following syntax is used to select an individual field from a record:

```
Variable#RecordName.Field
```

Note:

The values contained in record names and fields must be constants, not variables.

Note:

For the purposes of illustration, we will demonstrate the use of records using an imaginary dialogue with the Erlang shell. Currently the Erlang evaluator does not support records so you may not be able to reproduce this dialogue.

```
> P = #person{name = "Joe", phone = [0,8,2,3,4,3,1,2]}.  
{person, "Joe", [0,8,2,3,4,3,1,2], undefined}  
> P#person.name.  
"Joe"
```

Note:

Selectors for records are allowed in guards.

Updating a Record

The following syntax is used to create a new copy of the record with some of the fields changed. Only the fields to be changed need to be referred to, all other fields retain their old values.

```
OldVariable#RecordName{Field1 = NewValue1,  
                        ...,  
                        FieldM = NewValueM}
```

For example:

```
> P1 = #person{name="Joe", phone=[1,2,3], address="A street".  
{person, "Joe", [1,2,3], "A street"}  
> P2 = P1#person{name="Robert".  
{person, "Robert", [1,2,3], "A street"}
```

Type Testing

The following guard test is used to test the type of a record:

```
record(Variable, RecordName)
```

The following example shows that the guard succeeds if P is record of type person.

```
foo(P) when record(P, person) -> a_person;
foo(_) -> not_a_person.
```

Note:

This test checks that P is a tuple of arity $N + 1$, where N is the number of fields in the record, and the first element in the tuple is the atom person.

Pattern Matching

Matching can be used in combination with records as shown in the following example:

```
> P = #person{name="Joe", phone=[0,0,7], address="A street"}.
{person, "Joe", [0,0,7], "A street"}
> #person{name = Name} = P, Name.
"Joe"
```

The following function takes a list of person records and searches for the phone number of a person with a particular name:

```
find_phone([#person{name=Name, phone=Phone} | _], Name) ->
    {found, Phone};
find_phone([_| T], Name) ->
    find_phone(T, Name);
find_phone([], Name) ->
    not_found.
```

Note:

The fields referred to in the pattern can be given in any order.

Nested Records

The value of a field in a record might be an instance of a record. Retrieval of nested data can be done stepwise, or in a single step, as shown in the following example:

```
-record(person, {name = #name{ }, phone}).
-record(name, {first = "Robert", last = "Ericsson"}).

demo() ->
  P = #person{name= #name{first="Robert",last="Virding"}, phone=123},
  First = (P#person.name)#name.first.
```

Note:

In this example, `demo()` evaluates to "Robert".

Internal Representation of Records

It is often desirable to write generic functions which will work on any record, not just a record of a particular type. For this reason, records are represented internally as tuples and the ordering of the fields in the tuple is strictly defined.

For example, the record `-record(person, {name, phone, address}).` is represented internally by the tuple `{person, X, Y, Z}`.

The arity of the tuple is one more than the number of fields in the record. The first element of the tuple is the name of the record, and the elements of the tuple are the fields in the record. The variables `X`, `Y` and `Z` will store the data contained in the record fields.

The following two functions determine the indices in the tuple which refer to the named fields in the record:

- `record_info(fields, Rec) -> [Names]`. This function returns the names of the fields in the record `Rec`. For example, `record_info(fields, person)` evaluates to `[name, address, phone]`.
- `record_info(size, Rec) -> Size`. This function returns the size of the record `Rec` when represented as a tuple, which is one more than the number of fields. For example, `record_info(size, person)` returns 4.

In addition, `#Rec.Name` returns the index in the tuple representation of `Name` of the record `Rec`.

Note:

`Name` must be an atom.

For example, the following test function `test()` might return the result shown:

```
test() ->
    {record_info(fields, person),
      record_info(size, person),
      #person.name}.
```

```
> Mod:test().
{[name,address,phone],4,2}
```

The order in which records map onto tuples is implementation dependent.

Note:

`record_info` is a pseudo-function which cannot be exported from the module where it occurs.

Example

```
%% File: person.hrl

%%-----
%% Data Type: person
%% where:
%%   name:  A string (default is undefined).
%%   age:   An integer (default is undefined).
%%   phone: A list of integers (default is []).
%%   dict:  A dictionary containing various information
%%          about the person.
%%          A {Key, Value} list (default is the empty list).
%%-----
-record(person, {name, age, phone = [], dict = []}).

-module(person).
-include("person.hrl").
-compile(export_all). % For test purposes only.

%% This creates an instance of a person.
%% Note: The phone number is not supplied so the
%%       default value [] will be used.

make_hacker_without_phone(Name, Age) ->
    #person{name = Name, age = Age,
            dict = [{computer_knowledge, excellent},
                    {drinks, coke}]}.

%% This demonstrates matching in arguments

print(#person{name = Name, age = Age,
              phone = Phone, dict = Dict}) ->
    io:format("Name: ~s, Age: ~w, Phone: ~w ~n"
             "Dictionary: ~w.~n", [Name, Age, Phone, Dict]).
```

```
%% Demonstrates type testing, selector, updating.

birthday(P) when record(P, person) ->
    P#person{age = P#person.age + 1}.

register_two_hackers() ->
    Hacker1 = make_hacker_without_phone("Joe", 29),
    OldHacker = birthday(Hacker1),
    % The central_register_server should have
    % an interface function for this.
    central_register_server ! {register_person, Hacker1},
    central_register_server ! {register_person,
        OldHacker#person{name = "Robert",
            phone = [0,8,3,2,4,5,3,1]}}.
```

1.2 Programming with Funs

This section introduces functional objects (Funs), which are a new data type introduced in Erlang 4.4. Functions which takes Funs as arguments, or which return Funs are called higher order functions.

- Funs can be passed as arguments to other functions, just like lists or tuples
- functions can be written which return Funs, just like any other data object.

Higher Order Functions

Funs encourages us to encapsulate common patterns of design into functional forms called higher order functions. These functions not only shortens programs, but also produce clearer programs because the intended meaning of the program is explicitly rather than implicitly stated.

The concepts of higher order functions and procedural abstraction are introduced with two brief examples.

Example 1 - map

If we want to double every element in a list, we could write a function named `double`:

```
double([H|T]) -> [2*H|double(T)];
double([])    -> []
```

This function obviously doubles the argument entered as input as follows:

```
> double([1,2,3,4]).
[2,4,6,8]
```

We now add the function `add_one`, which adds one to every element in a list:

```
add_one([H|T]) -> [H+1|add_one(T)];
add_one([])    -> [] .
```

These functions, `double` and `add_one`, have a very similar structure. We can exploit this fact and write a function `map` which expresses this similarity:

```
map(F, [H|T]) -> [F(H)|map(F, T)];
map(F, [])    -> [] .
```

We can now express the functions `double` and `add_one` in terms of `map` as follows:

```
double(L) -> map(fun(X) -> 2*X end, L).
add_one(L) -> map(fun(X) -> 1 + X end, L).
```

`map(F, List)` is a function which takes a function `F` and a list `L` as arguments and returns the new list which is obtained by applying `F` to each of the elements in `L`.

The process of abstracting out the common features of a number of different programs is called procedural abstraction. Procedural abstraction can be used in order to write several different functions which have a similar structure, but differ only in some minor detail. This is done as follows:

1. write one function which represents the common features of these functions
2. parameterize the difference in terms of functions which are passed as arguments to the common function.

Example 2 - foreach

This example illustrates procedural abstraction. Initially, we show the following two examples written as conventional functions:

1. all elements of a list are printed onto a stream
2. a message is broadcast to a list of processes.

```
print_list(Stream, [H|T]) ->
  io:format(Stream, "~p~n", [H]),
  print_list(Stream, T);
print_on_list(Stream, []) ->
  true.
```

```
broadcast(Msg, [Pid|Pids]) ->
  Pid ! Msg,
  broadcast(Msg, Pids);
broadcast(_, []) ->
  true.
```

Both these functions have a very similar structure. They both iterate over a list doing something to each element in the list. The “something” has to be carried round as an extra argument to the function which does this.

The function `foreach` expresses this similarity:

```
foreach(F, [H|T]) ->
  F(H),
  foreach(F, T);
foreach(F, []) ->
  ok.
```

Using `foreach`, `print_on_list` becomes:

```
foreach(fun(H) -> io:format(S, "~p~n", [H]) end, L)
```

`broadcast` becomes:

```
foreach(fun(Pid) -> Pid ! M end, L)
```

`foreach` is evaluated for its side-effect and not its value. `foreach(Fun, L)` calls `Fun(X)` for each element `X` in `L` and the processing occurs in the order in which the elements were defined in `L`. `map` does not define the order in which its elements are processed.

Advantages of Higher Order Functions

Programming with higher order functions, such as `map` and `foreach`, has a number of advantages:

- It is much easier to understand the program and the intention of the programmer is clearly expressed in the code. The statement `foreach(fun(X) ->` clearly indicates that the intention of this program is to do something to each element in the list `L`. We also know that the function which is passed as the first argument of `foreach` takes one argument `X`, which will be successively bound to each of the elements in `L`.
- Functions which take Funs as arguments are much easier to re-use than other functions.

The Syntax of Funs

Funs are written with the syntax:

```
F = fun (Arg1, Arg2, ... ArgN) ->
      ...
      end
```

This creates an anonymous function of `N` arguments and binds it to the variable `F`.

If we have already written a function in the same module and wish to pass this function as an argument, we can use the following syntax:

```
F = fun FunctionName/Arity
```

With this form of function reference, the function which is referred to does not need to be exported from the module.

We can also refer to a function defined in a different module with the following syntax:

```
F = {Module, FunctionName}
```

In this case, the function must be exported from the module in question.

The follow program illustrates the different ways of creating Funs:

```
-module(fun_test).
-export([t1/0, t2/0, t3/0, t4/0, double/1]).
-import(lists, [map/2]).

t1() -> map(fun(X) -> 2 * X end, [1,2,3,4,5]).

t2() -> map(fun double/1, [1,2,3,4,5]).

t3() -> map({?MODULE, double}, [1,2,3,4,5]).

double(X) -> X * 2.
```

We can evaluate the fun `F` with the syntax:

```
F(Arg1, Arg2, ..., ArgN)
```

Variable Bindings within a Fun

The scope rules for variables which occur in Funs are as follows:

- All variables which occur in the head of a Fun are assumed to be “fresh” variables.
- Variables which are defined before the Fun, and which occur in function calls or guard tests within the Fun, have the values they had outside the Fun.
- No variables may be exported from a Fun.

The following examples illustrate these rules:

```
print_list(File, List) ->
  {ok, Stream} = file:open(File, write),
  foreach(fun(X) -> io:format(Stream, "~p~n", [X]) end, List),
  file:close(Stream).
```

In the above example, the variable `X` which is defined in the head of the Fun is a new variable. The value of the variable `Stream` which is used within within the Fun gets its value from the `file:open` line. Since any variable which occurs in the head of a Fun is considered a new variable it would be equally valid to write:

```
print_list(File, List) ->
  {ok, Stream} = file:open(File, write),
  foreach(fun(File) ->
    io:format(Stream, "~p~n", [File])
    end, List),
  file:close(Stream).
```

In this example, `File` is used as the new variable instead of `X`. This is rather silly since code in the body of the Fun cannot refer to the variable `File` which is defined outside the Fun. Compiling this example will yield the diagnostic:

```
./FileName.erl:Line: Warning: variable 'File'
  shadowed in 'lambda head'
```

This reminds us that the variable `File` which is defined inside the Fun collides with the variable `File` which is defined outside the Fun.

The rules for importing variables into a Fun has the consequence that certain pattern matching operations have to be moved into guard expressions and cannot be written in the head of the Fun. For example, we might write the following code if we intend the first clause of `F` to be evaluated when the value of its argument is `Y`:

```
f(...) ->
  Y = ...
  map(fun(X) when X == Y ->
    ;
    (_) ->
    ...
  end, ...)
```

instead of

```
f(...) ->
  Y = ...
  map(fun(Y) ->
      ;
      (_) ->
      ...
      end, ...)
  ...
```

Funs and the Module lists

The following examples show a dialogue with the Erlang shell. All the higher order functions discussed are exported from the module `lists`.

map

```
map(F, [H|T]) -> [F(H)|map(F, T)];
map(F, []) -> [].
```

`map` takes a function of one argument and a list of terms. It returns the list obtained by applying the function to every argument in the list.

```
1> Double = fun(X) -> 2 * X end.
#Fun<erl_eval>
2> lists:map(Double, [1,2,3,4,5]).
[2,4,6,8,10]
```

When a new Fun is defined in the shell, the value of the Fun is printed as `Fun#<erl_eval>`

any

```
any(Pred, [H|T]) ->
  case Pred(H) of
    true -> true;
    false -> any(Pred, T)
  end;
any(Pred, []) ->
  false.
```

`any` takes a predicate `P` of one argument and a list of terms. A predicate is a function which returns `true` or `false`. `any` is true if there is a term `X` in the list such that `P(X)` is true.

We define a predicate `Big(X)` which is true if its argument is greater than 10.

```
3> Big = fun(X) -> if X > 10 -> true; true -> false end end.
#Fun<erl_eval>
4> lists:any(Big, [1,2,3,4]).
false.
5> lists:any(Big, [1,2,3,12,5]).
true.
```

all

```
all(Pred, [H|T]) ->
  case Pred(H) of
    true -> all(Pred, T);
    false -> false
  end;
all(Pred, []) ->
  true.
```

all has the same arguments as any. It is true if the predicate applied to all elements in the list is true.

```
6> lists:all(Big, [1,2,3,4,12,6]).
false
7> lists:all(Big, [12,13,14,15]).
true
```

foreach

```
foreach(F, [H|T]) ->
  F(H),
  foreach(F, T);
foreach(F, []) ->
  ok.
```

foreach takes a function of one argument and a list of terms. The function is applied to each argument in the list. foreach returns ok. It is used for its side-effect only.

```
8> lists:foreach(fun(X) -> io:format("~w~n",[X]) end, [1,2,3,4]).
1
2
3
4
true
```

foldl

```
foldl(F, Accu, [Hd|Tail]) ->
  foldl(F, F(Hd, Accu), Tail);
foldl(F, Accu, []) -> Accu.
```

foldl takes a function of two arguments, an accumulator and a list. The function is called with two arguments. The first argument is the successive elements in the list, the second argument is the accumulator. The function must return a new accumulator which is used the next time the function is called.

If we have a list of lists L = ["I", "like", "Erlang"], then we can sum the lengths of all the strings in L as follows:

```
9> L = ["I", "like", "Erlang"].
["I", "like", "Erlang"]
10> lists:foldl(fun(X, Sum) -> length(X) + Sum end, 0, L).
11
```

foldl works like a while loop in an imperative language:

```
L = ["I","like","Erlang"],
Sum = 0,
while( L != []){
    Sum += length(head(L)),
    L = tail(L)
end
```

mapfoldl

```
mapfoldl(F, Accu0, [Hd|Tail]) ->
    {R,Accu1} = F(Hd, Accu0),
    {Rs,Accu2} = mapfoldl(F, Accu1, Tail),
    {[R|Rs], Accu2};
mapfoldl(F, Accu, []) -> {[], Accu}.
```

mapfoldl simultaneously maps and folds over a list. The following example shows how to change all letters in L to upper case and count them.

First upcase:

```
11> Upcase = fun(X) when $a =< X, X =< $z -> X + $A - $a;
      (X) -> X
      end.
#Fun<erl_eval>
12> Upcase_word =
      fun(X) ->
      lists:map(Upcase, X)
      end.
#Fun<erl_eval>
13> Upcase_word("Erlang").
"ERLANG"
14> lists:map(Upcase_word, L).
["I","LIKE","ERLANG"]
```

Now we can do the fold and the map at the same time:

```
14> lists:mapfoldl(fun(Word, Sum) ->
14> {Upcase_word(Word), Sum + length(Word)}
14> end, 0, L).
{"I","LIKE","ERLANG",11}
```

filter

```
filter(F, [H|T]) ->
  case F(H) of
    true  -> [H|filter(F, T)];
    false -> filter(F, T)
  end;
filter(F, []) -> [].
```

`filter` takes a predicate of one argument and a list and returns all element in the list which satisfy the predicate.

```
15> lists:filter(Big, [500,12,2,45,6,7]).
[500,12,45]
```

When we combine maps and filters we can write very succinct and obviously correct code. For example, suppose we want to define a set difference function. We want to define `diff(L1, L2)` to be the difference between the lists `L1` and `L2`. This is the list of all elements in `L1` which are not contained in `L2`. This code can be written as follows:

```
diff(L1, L2) ->
  filter(fun(X) -> not member(X, L2) end, L1).
```

The AND intersection of the list `L1` and `L2` is also easily defined:

```
intersection(L1,L2) -> filter(fun(X) -> member(X,L1) end, L2).
```

takewhile

```
takewhile(Pred, [H|T]) ->
  case Pred(H) of
    true  -> [H|takewhile(Pred, T)];
    false -> []
  end;
takewhile(Pred, []) ->
  [].
```

`takewhile(P, L)` takes elements `X` from a list `L` as long as the predicate `P(X)` is true.

```
16> lists:takewhile(Big, [200,500,45,5,3,45,6]).
[200,500,45]
```

dropwhile

```
dropwhile(Pred, [H|T]) ->
  case Pred(H) of
    true  -> dropwhile(Pred, T);
    false -> [H|T]
  end;
dropwhile(Pred, []) ->
  [].
```

dropwhile is the complement of takewhile.

```
17> lists:dropwhile(Big, [200,500,45,5,3,45,6]).
[5,3,45,6]
```

splitlist

```
splitlist(Pred, L) ->
  splitlist(Pred, L, []).

splitlist(Pred, [H|T], L) ->
  case Pred(H) of
    true  -> splitlist(Pred, T, [H|L]);
    false -> {reverse(L), [H|T]}
  end;
splitlist(Pred, [], L) ->
  {reverse(L), []}.
```

splitlist(P, L) splits the list L into the two sub-lists {L1, L2}, where L1 = takewhile(P, L) and L2 = dropwhile(P, L).

```
18> lists:splitlist(Big, [200,500,45,5,3,45,6]).
{[200,500,45],[5,3,45,6]}
```

first

```
first(Pred, [H|T]) ->
  case Pred(H) of
    true  ->
      {true, H};
    false ->
      first(Pred, T)
  end;
first(Pred, []) ->
  false.
```

first returns {true, R}, where R is the first element in a list satisfying a predicate or false:

```
19> lists:first(Big, [1,2,45,6,123]).
{true,45}
20> lists:first(Big, [1,2,4,5]).
false
```

Funs which Return Funs

So far, this section has only described functions which take Funs as arguments. It is also possible to write more powerful functions which themselves return Funs. The following examples illustrate these type of functions.

Simple Higher Order Functions

Adder(X) is a function which, given X, returns a new function G such that G(K) returns $K + X$.

```
21> Adder = fun(X) -> fun(Y) -> X + Y end end.
#Fun<erl.eval>
22> Add6 = Adder(6).
#Fun<erl.eval>
23> Add6(10).
16
```

Infinite Lists

The idea is to write something like:

```
-module(lazy).
-export([ints_from/1]).
ints_from(N) ->
    fun() ->
        [N|ints_from(N+1)]
    end.
```

Then we can proceed as follows:

```
24> XX = lazy:ints_from(1).
#Fun<lazy>
25> XX().
[1|#Fun<lazy>]
26> hd(XX()).
1
27> Y = tl(XX()).
#Fun<lazy>
28> hd(Y()).
2
```

etc. - this is an example of “lazy embedding”

Parsing

The following examples show parsers of the following type:

```
Parser(Toks) -> {ok, Tree, Toks1} | fail
```

Toks is the list of tokens to be parsed. A successful parse returns {ok, Tree, Toks1}, where Tree is a parse tree and Toks1 is a tail of Tree which contains symbols encountered after the structure which was correctly parsed. Otherwise fail is returned.

The example which follows illustrates a simple, functional parser which parses the grammar:

```
(a | b) & (c | d)
```

The following code defines a function pconst(X) in the module funparse, which returns a Fun which parses a list of tokens.

```
pconst(X) ->
  fun (T) ->
    case T of
      [X|T1] -> {ok, {const, X}, T1};
      _      -> fail
    end
  end.
```

This function can be used as follows:

```
29> P1 = funparse:pconst(a).
#Fun<hof>
30> P1([a,b,c]).
{ok, {const, a}, [b,c]}
31> P1([x,y,z]).
fail
```

Next, we define the two higher order functions pand and por which combine primitive parsers to produce more complex parsers. Firstly pand:

```
pand(P1, P2) ->
  fun (T) ->
    case P1(T) of
      {ok, R1, T1} ->
        case P2(T1) of
          {ok, R2, T2} ->
            {ok, {'and', R1, R2}};
          fail ->
            fail
        end;
      fail ->
        fail
    end
  end.
```

Given a parser P1 for grammar G1, and a parser P2 for grammar G2, `pand(P1, P2)` returns a parser for the grammar which consists of sequences of tokens which satisfy G1 followed by sequences of tokens which satisfy G2.

`por(P1, P2)` returns a parser for the language described by the grammar G1 or G2.

```
por(P1, P2) ->
  fun (T) ->
    case P1(T) of
      {ok, R, T1} ->
        {ok, {'or',1,R}, T1};
      fail ->
        case P2(T) of
          {ok, R1, T1} ->
            {ok, {'or',2,R1}, T1};
          fail ->
            fail
        end
    end
  end
end.
```

The original problem was to parse the grammar $(a \mid b) \& (c \mid d)$. The following code addresses this problem:

```
grammar() ->
  pand(
    por(pconst(a), pconst(b)),
    por(pconst(c), pconst(d))).
```

The following code adds a parser interface to the grammar:

```
parse(List) ->
  (grammar())(List).
```

We can test this parser as follows:

```
32> funparse:parse([a,c]).
{ok,{'and',{'or',1,{const,a}},{'or',1,{const,c}}}}
33> funparse:parse([a,d]).
{ok,{'and',{'or',1,{const,a}},{'or',2,{const,d}}}}
34> funparse:parse([b,c]).
{ok,{'and',{'or',2,{const,b}},{'or',1,{const,c}}}}
35> funparse:parse([b,d]).
{ok,{'and',{'or',2,{const,b}},{'or',2,{const,d}}}}
36> funparse:parse([a,b]).
fail
```

1.3 List Comprehensions

List comprehensions are a feature of many modern functional programming languages. Subject to certain rules, they provide a succinct notation for generating elements in a list.

List comprehensions are analogous to set comprehensions in Zermelo-Frankel set theory and are called ZF expressions in Miranda and SASL. They are analogous to the `setof` and `findall` predicates in Prolog.

List comprehensions are written with the following syntax:

```
[Expression || Qualifier1, Qualifier2, ...]
```

Expression is an arbitrary expression, and each Qualifier is either a generator or a filter.

- A *generator* written as `Pattern <- ListExpr`. `ListExpr` must be an expression which evaluates to a list of terms.
- A *filter* is either a predicate or a boolean expression. A predicate is a function which returns `true` or `false`.

Examples of List Comprehensions

We start with a simple example:

```
> [X || X <- [1,2,a,3,4,b,5,6], X > 3].  
[a,4,b,5,6]
```

This should be read as follows:

The list of X such that X is taken from the list `[1,2,a,...]` and X is greater than 3.

The notation `X <- [1,2,a,...]` is a generator and the expression `X > 3` is a filter.

An additional filter can be added in order to restrict the result to integers:

```
> [X || X <- [1,2,a,3,4,b,5,6], integer(X), X > 3].  
[4,5,6]
```

Generators can be combined. For example, the Cartesian product of two lists can be written as follows:

```
> [{X, Y} || X <- [1,2,3], Y <- [a,b]].  
[{1,a},{1,b},{2,a},{2,b},{3,a},{3,b}]
```

Quick Sort

The well known quick sort routine can be written as follows:

```
sort([Pivot|T]) ->
  sort([ X || X <- T, X < Pivot]) ++
  [Pivot] ++
  sort([ X || X <- T, X >= Pivot]);
sort([]) -> [].
```

The expression `[X || X <- T, X < Pivot]` is the list of all elements in `T`, which are less than `Pivot`.

`[X || X <- T, X >= Pivot]` is the list of all elements in `T`, which are greater or equal to `Pivot`.

To sort a list, we isolate the first element in the list and split the list into two sub-lists. The first sub-list contains all elements which are smaller than the first element in the list, the second contains all elements which are greater than or equal to the first element in the list. We then sort the sub-lists and combine the results.

Permutations

The following example generates all permutations of the elements in a list:

```
perms([]) -> [[]];
perms(L) -> [[H|T] || H <- L, T <- perms(L--[H])].
```

We take `H` from `L` in all possible ways. The result is the set of all lists `[H|T]`, where `T` is the set of all possible permutations of `L` with `H` removed.

```
> perms([b,u,g]).
[[b,u,g],[b,g,u],[u,b,g],[u,g,b],[g,b,u],[g,u,b]]
```

Pythagorean Triplets

Pythagorean triplets are sets of integers $\{A,B,C\}$ such that $A^2 + B^2 = C^2$.

The function `pyth(N)` generates a list of all integers $\{A,B,C\}$ such that $A^2 + B^2 = C^2$ and where the sum of the sides is less than `N`.

```
pyth(N) ->
  [ {A,B,C} ||
    A <- lists:seq(1,N),
    B <- lists:seq(1,N),
    C <- lists:seq(1,N),
    A+B+C =< N,
    A*A+B*B == C*C
  ].
```

```

> pyth(3).
[] .
> pyth(11).
[] .
> pyth(12).
[{3,4,5},{4,3,5}]
> pyth(50).
[{3,4,5},{4,3,5},{5,12,13},{6,8,10},{8,6,10},{8,15,17},
{9,12,15},{12,5,13},{12,9,15},{12,16,20},{15,8,17},
{16,12,20}]

```

The following code reduces the search space and is more efficient:

```

pyth1(N) ->
  [{A,B,C} ||
    A <- lists:seq(1,N),
    B <- lists:seq(1,N-A+1),
    C <- lists:seq(1,N-A-B+2),
    A+B+C =< N,
    A*A+B*B == C*C ].

```

Simplifications with List Comprehensions

As an example, list comprehensions can be used to simplify some of the functions in `lists.erl`:

```

append(L)  -> [X || L1 <- L, X <- L1].
map(Fun, L) -> [Fun(X) || X <- L].
filter(Pred, L) -> [X || X <- L, Pred(X)].

```

Variable Bindings in List Comprehensions

The scope rules for variables which occur in list comprehensions are as follows:

- all variables which occur in a generator pattern are assumed to be “fresh” variables
- any variables which are defined before the list comprehension and which are used in filters have the values they had before the list comprehension
- no variables may be exported from a list comprehension.

As an example of these rules, suppose we want to write the function `select`, which selects certain elements from a list of tuples. We might write `select(X, L) -> [Y || {X, Y} <- L]` with the intention of extracting all tuples from `L` where the first item is `X`.

Compiling this yields the following diagnostic:

```
./FileName.erl:Line: Warning: variable 'X' shadowed in generate
```

This diagnostic warns us that the variable `X` in the pattern is not the same variable as the variable `X` which occurs in the function head.

Evaluating `select` yields the following result:

```
> select(b, [{a,1},{b,2},{c,3},{b,7}]).  
[1,2,3,7]
```

This result is not what we wanted. To achieve the desired effect we must write `select` as follows:

```
select(X, L) -> [Y || {X1, Y} <- L, X == X1].
```

The generator now contains unbound variables and the test has been moved into the filter. This now works as expected:

```
> select(b, [{a,1},{b,2},{c,3},{b,7}]).  
[2,7]
```

One consequence of the rules for importing variables into a list comprehensions is that certain pattern matching operations have to be moved into the filters and cannot be written directly in the generators. To illustrate this, do *not* write as follows:

```
f(...) ->  
  Y = ...  
  [ Expression || PatternInvolving Y <- Expr, ...]  
  ...
```

Instead, write as follows:

```
f(...) ->  
  Y = ...  
  [ Expression || PatternInvolving Y1 <- Expr, Y == Y1, ...]  
  ...
```

1.4 Macros

Macros in Erlang are written with the following syntax:

```
-define(Const, Replacement).  
-define(Fun(Var1, Var2, ..., Var), Replacement).
```

Macros are expanded when the syntax `?MacroName` is encountered.

Consider the macro definition:

```
-define(timeout, 200).
```

The expression `?timeout`, which can occur anywhere in the code which follows the macro definition, will be replaced by `200`.

Macros with arguments are written as follows:

```
-define(macro1(X, Y), {a, X, b, Y}).
```

This type of macro can be used as follows:

```
bar(X) ->  
    ?macro1(a, b),  
    ?macro1(X, 123)
```

This expands to:

```
bar(X) ->  
    {a, a, b, b},  
    {a, X, b, 123}.
```

Macros and Tokens

Macro expansion works at a token level. We might define a macro as follows:

```
-define(macro2(X, Y), {a,X,b,Y}).
```

The replacement value of the macro is not a valid Erlang term because the closing right curly bracket is missing. `macro2` expands into a sequence of tokens `{, a, X` which are then pasted into the place where the macro is used.

We might use this macro as follows:

```
bar() ->  
    ?macro2(x,y)}.
```

This will expand into the valid sequence of tokens {a, x, y, b} before being parsed and compiled.

Note:

It is good programming practise to ensure that the replacement text of a macro is a valid Erlang syntactic form.

Pre-Defined Macros

The following macros are pre-defined:

?MODULE. This macro returns the name of the current module

?FILE. This macro returns the current file name.

?LINE. This macro returns the current line number.

?MACHINE. This macro returns the current machine name, which is one of 'JAM', 'BEAM' or 'VEE'

Flow Control in Macros

The following macro directives are supplied:

-undef(Macro). Causes the macro to behave as if it had never been defined.

-ifdef(Macro). Do the following lines if Macro is defined.

-ifndef(Macro). Do the following lines if Macro is not defined.

-else. "else" macro

-endif. "endif" macro.

The conditional macros must be properly nested. They are usually grouped as follows:

```
-ifdef(debug)
-define(...)
-else
-define(...)
-endif
```

The following example illustrates this grouping:

```
-define(debug, true).
-ifdef(debug).
-define(trace(Str, X), io:format("Mod:~w line:~w ~p ~p~n",
                                [?MODULE,?LINE, Str, X])).
-else.
-define(trace(X, Y), true).
-endif.
```

Given these definitions, the expression ?trace("X=", X). in line 10 of the module foo expands to:

```
io:format("Mod:~w line:~w ~p ~p~n",[foo,100,"X=", [X]]),
```

If we remove the `-define(debug, true).` line, then the same expression expands to `true`.

A Macro Expansion Utility

The following code can be used to expand a macro and display the result:

```
-module(mexpand).
-export([file/1]).
-import(lists, [foreach/2]).
file(File) ->
  case epp:parse_file(File ++ ".erl", [], []) of
    {ok, L} ->
      {ok, Stream} = file:open(File ++ ".out", write),
      foreach(fun(X) ->
                io:format(Stream, "~s~n", [erl_pp:form(X)])
              end, L),
      file:close(Stream)
  end.
```

Alternatively, we can compile the file with the `'P'` option. `compile:file(File, ['P'])` produces a list file `File.P`, in which the result of any macro expansions can be seen.